Published before May 2001
Terms used **bandwidth symmetric multiprocessor interconnect**

Found **48** of **113,585**

Sort results by: relevance

Display results: expanded form

Save results to a Binder

Search Tips

Open results in a new window

Try an Advanced Search

Try this search in The ACM Guide

Results 1 - 20 of 48 Result page: **1** 2 3 next

Relevance scale

1 The Starfire SMP interconnect

Alan Charlesworth, Nicholas Aneshansley, Mark Haakmeester, Dan Drogichen, Gary Gilbert, Ricki Williams, Andrew Phelps

November 1997 **Proceedings of the 1997 ACM/IEEE conference on Supercomputing (CDROM)**

Full text available: pdf(273.52 KB) Additional Information: full citation, abstract, references, citings

The Starfire interconnect extends the envelope of Unix symmetric multiprocessor (SMP) systems in several dimensions. **Interconnect:** an active centerplane with four address routers and a 16x16 data crossbar provides 64 UltraSPARC processors with uniform memory access at a bandwidth of 10,667 MBps. **Flexibility:** Starfire can be dynamically reconfigured into multiple hardware-protected operating system domains. **Robustness:** Failing boards can be hot swapped without interrupting sy ...

Keywords: SMP, UMA, bandwidth, domains, interconnect, latency, partitions

2 Accelerating shared virtual memory via general-purpose network interface support

Angelos Bilas, Dongming Jiang, Jaswinder Pal Singh

February 2001 **ACM Transactions on Computer Systems (TOCS)**, Volume 19 Issue 1

Full text available: pdf(178.88 KB) Additional Information: full citation, abstract, references, index terms, review

Clusters of symmetric multiprocessors (SMPs) are important platforms for high-performance computing. With the success of hardware cache-coherent distributed shared memory (DSM), a lot of effort has also been made to support the coherent shared-address-space programming model in software on clusters. Much research has been done in fast communication on clusters and in protocols for supporting software shared memory across them. However, the performance of software virtual memory (SVM) is sti ...

Keywords: applications, clusters, shared virtual memory, system area networks

3 Multi-protocol active messages on a cluster of SMP's

Steven S. Lumetta, Alan M. Mainwaring, David E. Culler

November 1997 **Proceedings of the 1997 ACM/IEEE conference on Supercomputing**

(CDROM)

Full text available: pdf(248.27 KB) Additional Information: full citation, abstract, references, citings

Clusters of multiprocessors, or Clumps, promise to be the supercomputers of the future, but obtaining high performance on these architectures requires an understanding of interactions between the multiple levels of interconnection. In this paper, we present the first multi-protocol implementation of a lightweight message layer---a version of Active Messages-II running on a cluster of Sun Enterprise 5000 servers connected with Myrinet. This research brings together several pieces of high-performa ...

4 FM-QoS: real-time communication using self-synchronizing schedules

Kay Connelly, Andrew A. Chien

November 1997 **Proceedings of the 1997 ACM/IEEE conference on Supercomputing (CDROM)**

Full text available: pdf(145.06 KB) Additional Information: full citation, abstract, references, citings

FM-QoS employs a novel communication architecture based on network feedback to provide predictable communication performance (e.g. deterministic latencies and guaranteed bandwidths) for high speed cluster interconnects. Network feedback is combined with self-synchronizing communication schedules to achieve synchrony in the network interfaces (NIs). Based on this synchrony, the network can be scheduled to provide predictable performance without special network QoS hardware. We describe the key el ...

Keywords: communication, network, predictable performance, quality-of-service, real-time, scheduling, synchronization, wormhole

5 The SGI Origin: a ccNUMA highly scalable server

James Laudon, Daniel Lenoski

May 1997 **ACM SIGARCH Computer Architecture News , Proceedings of the 24th annual international symposium on Computer architecture**, Volume 25 Issue 2

Full text available: pdf(1.74 MB) Additional Information: full citation, abstract, references, citings, index terms

The SGI Origin 2000 is a cache-coherent non-uniform memory access (ccNUMA) multiprocessor designed and manufactured by Silicon Graphics, Inc. The Origin system was designed from the ground up as a multiprocessor capable of scaling to both small and large processor counts without any bandwidth, latency, or cost cliffs. The Origin system consists of up to 512 nodes interconnected by a scalable Craylink network. Each node consists of one or two R10000 processors, up to 4 GB of coherent memory, and ...

6 Parallel implementation of a molecular dynamics simulation program

Alan Mink, Christophe Bailly

December 1998 **Proceedings of the 30th conference on Winter simulation**

Full text available: pdf(100.62 KB) Additional Information: full citation, references, index terms

7 Timestamp snooping: an approach for extending SMPs

Milo M. K. Martin, Daniel J. Sorin, Anastassia Ailamaki, Alaa R. Alameldeen, Ross M. Dickson, Carl J. Mauer, Kevin E. Moore, Manoj Plakal, Mark D. Hill, David H. Wood

November 2000 **ACM SIGPLAN Notices**, Volume 35 Issue 11

Full text available: pdf(1.30 MB) Additional Information: full citation, abstract, references, citings, index terms

Symmetric multiprocessor (SMP) servers provide superior performance for the commercial workloads that dominate the Internet. Our simulation results show that over one-third of

cache misses by these applications result in cache-to-cache transfers, where the data is found in another processor's cache rather than in memory. SMPs are optimized for this case by using snooping protocols that broadcast address transactions to all processors. Conversely, directory-based shared-memory systems must indire ...

8 STiNG: a CC-NUMA computer system for the commercial marketplace

Tom Lovett, Russell Clapp

May 1996 **ACM SIGARCH Computer Architecture News , Proceedings of the 23rd annual international symposium on Computer architecture**, Volume 24 Issue 2

Full text available: pdf(1.30 MB) Additional Information: full citation, abstract, references, citings, index terms

"STiNG" is a Cache Coherent Non-Uniform Memory Access (CC-NUMA) Multiprocessor designed and built by Sequent Computer Systems, Inc. It combines four processor Symmetric Multi-processor (SMP) nodes (called Quads), using a Scalable Coherent Interface (SCI) based coherent interconnect. The Quads are based on the Intel P6 processor and the external bus it defines. In addition to 4 P6 processors, each Quad may contain up to 4 GBytes of system memory, 2 Peripheral Component Interface (PCI) busses for ...

9 Timestamp snooping: an approach for extending SMPs

Milo M. K. Martin, Daniel J. Sorin, Anatassia Ailamaki, Alaa R. Alameldeen, Ross M. Dickson, Carl J. Mauer, Kevin E. Moore, Manoj Plakal, Mark D. Hill, David A. Wood

November 2000 **Proceedings of the ninth international conference on Architectural support for programming languages and operating systems**, Volume 28 , 34 Issue 5 , 5

Full text available: pdf(164.27 KB) Additional Information: full citation, abstract, references, citings, index terms

Symmetric muultiprocessor (SMP) servers provide superior performance for the commercial workloads that dominate the Internet. Our simulation results show that over one-third of cache misses by these applications result in cache-to-cache transfers, where the data is found in another processor's cache rather than in memory. SMPs are optimized for this case by using snooping protocols that broadcast address transactions to all processors. Conversely, directory-based shared-memory systems must indir ...

10 Micro-analysis of the titans's operating pipe

J. Sanguinetti

June 1988 **Proceedings of the 2nd international conference on Supercomputing**

Full text available: pdf(594.77 KB) Additional Information: full citation, abstract, references, citings, index terms

Much of the performance analysis done in designing a computer is based on fundamental operation rates, like cycle time and number of pipe stages in an operation pipeline. This kind of analysis yields peak computation rates which, in fact, may never be realized. Resource contention between different units, each of which has a fundamental operation rate adequate to support a given overall peak, may cause the actual obtainable rate to be much less. In order to determine the effects of interact ...

11 A personal supercomputer for climate research

James C. Hoe, Chris Hill, Alistair Adcroft

January 1999 **Proceedings of the 1999 ACM/IEEE conference on Supercomputing (CDROM)**

Full text available: pdf(491.63 KB) Additional Information: full citation, references, index terms

12 A case for intelligent disks (IDISKs)

Kimberly Keeton, David A. Patterson, Joseph M. Hellerstein
September 1998 **ACM SIGMOD Record**, Volume 27 Issue 3

Full text available: pdf(1.07 MB) Additional Information: full citation, abstract, citings, index terms

Decision support systems (DSS) and data warehousing workloads comprise an increasing fraction of the database market today. I/O capacity and associated processing requirements for DSS workloads are increasing at a rapid rate, doubling roughly every nine to twelve months [38]. In response to this increasing storage and computational demand, we present a computer architecture for decision support database servers that utilizes "intelligent" disks (IDISKs). IDISKs utilize low-cost ...

13 Performance evaluation of a commercial cache-coherent shared memory multiprocessor

Rajeev Jog, Philip L. Vitale, James R. Callister
April 1990 **ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 1990 ACM SIGMETRICS conference on Measurement and modeling of computer systems**, Volume 18 Issue 1

Full text available: pdf(948.46 KB) Additional Information: full citation, abstract, references, citings, index terms

This paper describes an approximate Mean Value Analysis (MVA) model developed to project the performance of a small-scale shared-memory commercial symmetric multiprocessor system. The system, based on Hewlett Packard Precision Architecture processors, supports multiple active user processes and multiple execution threads within the operating system. Using detailed timing for hardware delays, a customized approximate closed queueing model is developed for the multiprocessor system ...

14 Memory bandwidth limitations of future microprocessors

Doug Burger, James R. Goodman, Alain Kägi
May 1996 **ACM SIGARCH Computer Architecture News , Proceedings of the 23rd annual international symposium on Computer architecture**, Volume 24 Issue 2

Full text available: pdf(1.60 MB) Additional Information: full citation, abstract, references, citings, index terms

This paper makes the case that pin bandwidth will be a critical consideration for future microprocessors. We show that many of the techniques used to tolerate growing memory latencies do so at the expense of increased bandwidth requirements. Using a decomposition of execution time, we show that for modern processors that employ aggressive memory latency tolerance techniques, wasted cycles due to insufficient bandwidth generally exceed those due to raw memory latencies. Given the importance of ma ...

15 Can shared-memory model serve as a bridging model for parallel computation?

Phillip B. Gibbons, Yossi Matias, Vijaya Ramachandran
June 1997 **Proceedings of the ninth annual ACM symposium on Parallel algorithms and architectures**

Full text available: pdf(1.62 MB) Additional Information: full citation, references, citings, index terms

16 Architecture

Paul Messina, David Culler, Wayne Pfeiffer, William Martin, J. Tinsley Oden, Gary Smith
November 1998 **Communications of the ACM**, Volume 41 Issue 11

Full text available: pdf(334.29 KB) Additional Information: full citation, references, citings, index terms, review

17

Communication overlap in multi-tier parallel algorithms

Scott B. Baden, Stephen J. Fink

November 1998 **Proceedings of the 1998 ACM/IEEE conference on Supercomputing (CDROM)**

Full text available: pdf(278.73 KB) Additional Information: full citation, abstract, references

Hierarchically organized multicomputers such as SMP clusters offer new opportunities and new challenges for high-performance computation, but realizing their full potential remains a formidable task. We present a hierarchical model of communication targeted to block-structured, bulk-synchronous applications running on dedicated clusters of symmetric multiprocessors. Our model supports node-level rather processor-level communication as the fundamental operation, and is optimized for aggregate pat ...

18 Coherent network interfaces for fine-grain communication

Shubhendu S. Mukherjee, Babak Falsafi, Mark D. Hill, David A. Wood

May 1996 **ACM SIGARCH Computer Architecture News , Proceedings of the 23rd annual international symposium on Computer architecture**, Volume 24 Issue 2

Full text available: pdf(1.72 MB) Additional Information: full citation, abstract, references, citings, index terms

Historically, processor accesses to memory-mapped device registers have been marked uncachable to insure their visibility to the device. The ubiquity of snooping cache coherence, however, makes it possible for processors and devices to interact with cachable, coherent memory operations. Using coherence can improve performance by facilitating burst transfers of whole cache blocks and reducing control overheads (e.g., for polling).This paper begins an exploration of network interfaces (NIs) that u ...

19 Oracle media server: providing consumer based interactive access to multimedia data

Andrew Laursen, Jeffrey Olkin, Mark Porter

May 1994 **ACM SIGMOD Record , Proceedings of the 1994 ACM SIGMOD international conference on Management of data**, Volume 23 Issue 2

Full text available: pdf(1.05 MB) Additional Information: full citation, abstract, citings, index terms

Currently, most data accessed on large servers is structured data stored in traditional databases. Networks are LAN based and clients range from simple terminals to powerful workstations. The user is corporate and the application developer is an MIS professional.With the introduction of broadband communications to the home and better than 100-to-1 compression techniques, a new form of network-based computing is emerging. Structured data is still important, but the bulk of data b ...

20 Data prefetch mechanisms

Steven P. Vanderwiel, David J. Lilja

June 2000 **ACM Computing Surveys (CSUR)**, Volume 32 Issue 2

Full text available: pdf(172.07 KB) Additional Information: full citation, abstract, references, citings, index terms, review

The expanding gap between microprocessor and DRAM performance has necessitated the use of increasingly aggressive techniques designed to reduce or hide the latency of main memory access. Although large cache hierarchies have proven to be effective in reducing this latency for the most frequently used data, it is still not uncommon for many programs to spend more than half their run times stalled on memory requests. Data prefetching has been proposed as a technique for hiding the access lat ...

Keywords: memory latency, prefetching

Terms used **bandwidth symmetric multiprocessor interconnect alternate backup path link**

Found **2** of **150,138**

Sort results by: relevance

Display results: expanded form

Save results to a Binder

Search Tips

Open results in a new window

Try an Advanced Search

Try this search in The ACM Guide

Results 1 - 2 of 2

Relevance scale

1 A case for intelligent disks (IDISKs)

Kimberly Keeton, David A. Patterson, Joseph M. Hellerstein

September 1998 **ACM SIGMOD Record**, Volume 27 Issue 3

Full text available: pdf(1.07 MB) Additional Information: full citation, abstract, citings, index terms

Decision support systems (DSS) and data warehousing workloads comprise an increasing fraction of the database market today. I/O capacity and associated processing requirements for DSS workloads are increasing at a rapid rate, doubling roughly every nine to twelve months [38]. In response to this increasing storage and computational demand, we present a computer architecture for decision support database servers that utilizes "intelligent" disks (IDISKs). IDISKs utilize low-cost ...

2 Network attached storage architecture

Garth A. Gibson, Rodney Van Meter

November 2000 **Communications of the ACM**, Volume 43 Issue 11

Full text available: pdf(224.67 KB) html(43.39 KB) Additional Information: full citation, references, citings, index terms

Results 1 - 2 of 2

Nothing Found

Your search for **+" bandwidth" +"symmetric multiprocessor" +"backup interconnect"** did not return any results.

You may want to try an Advanced Search for additional options.

Please review the Quick Tips below or for more information see the Search Tips.

Quick Tips

- Enter your search terms in lower case with a space between the terms.

  ```
  sales offices
  ```

  You can also enter a full question or concept in plain language.

  ```
  Where are the sales offices?
  ```

- Capitalize proper nouns to search for specific people, places, or products.

  ```
  John Colter, Netscape Navigator
  ```

- Enclose a phrase in double quotes to search for that exact phrase.

  ```
  "museum of natural history" "museum of modern art"
  ```

- Narrow your searches by using a + if a search term must appear on a page.

  ```
  museum +art
  ```

- Exclude pages by using a - if a search term must not appear on a page.

  ```
  museum -Paris
  ```

Combine these techniques to create a specific search query. The better your description of the information you want, the more relevant your results will be.

```
museum +"natural history" dinosaur -Chicago
```

Nothing Found

Your search for **+" bandwidth" +"symmetric multiprocessor" +"backup connect"** did not return any results.

You may want to try an Advanced Search for additional options.

Please review the Quick Tips below or for more information see the Search Tips.

Quick Tips

- Enter your search terms in lower case with a space between the terms.

  ```
  sales offices
  ```

  You can also enter a full question or concept in plain language.

  ```
  Where are the sales offices?
  ```

- Capitalize proper nouns to search for specific people, places, or products.

  ```
  John Colter, Netscape Navigator
  ```

- Enclose a phrase in double quotes to search for that exact phrase.

  ```
  "museum of natural history" "museum of modern art"
  ```

- Narrow your searches by using a + if a search term must appear on a page.

  ```
  museum +art
  ```

- Exclude pages by using a - if a search term must not appear on a page.

  ```
  museum -Paris
  ```

Combine these techniques to create a specific search query. The better your description of the information you want, the more relevant your results will be.

```
museum +"natural history" dinosaur -Chicago
```

Nothing Found

Your search for **+"not sufficient bandwidth" +"multiprocessor"** did not return any results.

You may want to try an Advanced Search for additional options.

Please review the Quick Tips below or for more information see the Search Tips.

Quick Tips

- Enter your search terms in lower case with a space between the terms.

  ```
  sales offices
  ```

  You can also enter a full question or concept in plain language.

  ```
  Where are the sales offices?
  ```

- Capitalize proper nouns to search for specific people, places, or products.

  ```
  John Colter, Netscape Navigator
  ```

- Enclose a phrase in double quotes to search for that exact phrase.

  ```
  "museum of natural history" "museum of modern art"
  ```

- Narrow your searches by using a + if a search term must appear on a page.

  ```
  museum +art
  ```

- Exclude pages by using a - if a search term must not appear on a page.

  ```
  museum -Paris
  ```

Combine these techniques to create a specific search query. The better your description of the information you want, the more relevant your results will be.

```
museum +"natural history" dinosaur -Chicago
```

Published before May 2001
Terms used **not enough bandwidth multiprocessor**

Found **2** of **113,585**

Sort results by: relevance

Display results: expanded form

Save results to a Binder

Search Tips

Open results in a new window

Try an Advanced Search

Try this search in The ACM Guide

Results 1 - 2 of 2

Relevance scale

1 Protocol architectures: MMTP: multimedia multiplexing transport protocol

Luiz Magalhaes, Robin Kravets

April 2001 **ACM SIGCOMM Computer Communication Review**, Volume 31 Issue 2 supplement

Full text available: pdf(2.08 MB) Additional Information: full citation, abstract, references

Multimedia data has special requirements that are hard to be met on mobile hosts due to potentially low bandwidth and disruptions due to host mobility. Such limited communication capabilities of mobile hosts can be offset by the simultaneous use of multiple link layer technologies. MMTP is a member of a suite of protocols that share the novel characteristic of aggregating bandwidth from multiple link-layer channels. The use of multiple channels to transport user data provides five key benefits: ...

Keywords: low bandwidth link, multimedia transport protocols, wireless communication

2 A failure and overload tolerance mechanism for continuous media servers

Rajesh Krishnan, Dinesh Venkatesh, Thomas D. C. Little

November 1997 **Proceedings of the fifth ACM international conference on Multimedia**

Full text available: pdf(2.23 MB) Additional Information: full citation, references, index terms

Keywords: caching, clustered video servers, content insertion, fault tolerance, interactive video-on-demand, overload tolerance, rate adaptive stream merging, stream clustering

Results 1 - 2 of 2

Nothing Found

Your search for **+"calculate bandwidth" +"symmetric multiprocessor"** did not return any results.

You may want to try an Advanced Search for additional options.

Please review the Quick Tips below or for more information see the Search Tips.

Quick Tips

- Enter your search terms in lower case with a space between the terms.

  ```
  sales offices
  ```

  You can also enter a full question or concept in plain language.

  ```
  Where are the sales offices?
  ```

- Capitalize proper nouns to search for specific people, places, or products.

  ```
  John Colter, Netscape Navigator
  ```

- Enclose a phrase in double quotes to search for that exact phrase.

  ```
  "museum of natural history" "museum of modern art"
  ```

- Narrow your searches by using a + if a search term must appear on a page.

  ```
  museum +art
  ```

- Exclude pages by using a - if a search term must not appear on a page.

  ```
  museum -Paris
  ```

Combine these techniques to create a specific search query. The better your description of the information you want, the more relevant your results will be.

```
museum +"natural history" dinosaur -Chicago
```